Dictionary learning of sound speed profiles

Michael Bianco[a] and Peter Gerstoft
Scripps Institution of Oceanography, University of California San Diego, La Jolla, California 92093–0238

(Dated: March 17, 2017)

To provide constraints on their inversion, ocean sound speed profiles (SSPs) often are modeled using empirical orthogonal functions (EOFs). However, this regularization, which uses the leading order EOFs with a minimum-energy constraint on their coefficients, often yields low resolution SSP estimates. In this paper, it is shown that dictionary learning, a form of unsupervised machine learning, can improve SSP resolution by generating a dictionary of shape functions for sparse processing (e.g. compressive sensing) that optimally compress SSPs; both minimizing the reconstruction error and the number of coefficients. These learned dictionaries (LDs) are not constrained to be orthogonal and thus, fit the given signals such that each signal example is approximated using few LD entries. Here, LDs describing SSP observations from the HF-97 experiment and the South China Sea are generated using the K-SVD algorithm. These LDs better explain SSP variability and require fewer coefficients than EOFs, describing much of the variability with one coefficient. Thus, LDs improve the resolution of SSP estimates with negligible computational burden.
© 2016 Acoustical Society of America

Keywords: Ocean acoustics; geophysics; dictionary learning; machine learning; compressive sensing

PACS numbers: 43.60Pt, 43.30Pc

arXiv:1609.04840v3 [physics.ao-ph] 15 Mar 2017

I. INTRODUCTION

Inversion for ocean sound speed profiles (SSPs) using acoustic data is a non-linear and highly underdetermined problem.[1] To ensure physically realistic solutions while moderating the size of the parameter search, SSP inversion has often been regularized by modeling SSP as the sum of leading order empirical orthogonal functions (EOFs).[2–7] However, regularization using EOFs often yields low resolution estimates of ocean SSPs, which can be highly variable with fine scale fluctuations. In this paper, it is shown that the resolution of SSP estimates are improved using dictionary learning,[8–13] a form of unsupervised machine learning, to generate a dictionary of regularizing shape functions from SSP data for parsimonious representation of SSPs.

Many signals, including natural images[14,15], audio[16], and seismic profiles[17] are well approximated using sparse (few) coefficients, provided a dictionary of shape functions exist under which their representation is sparse. Given a K-dimensional signal, a dictionary is defined as a set of N, ℓ_2-normalized vectors which describe the signal using few coefficients. The sparse processor is then an ℓ_2-norm cost function with an ℓ_0-norm penalty on the number of non-zero coefficients. Signal sparsity is exploited for a number of purposes including signal compression and denoising.[9] Applications of compressive sensing,[18] one approximation to the ℓ_0-norm sparse processor, have in ocean acoustics shown improvements in beamforming,[19–22] geoacoustic inversion,[23] and estimation of ocean SSPs.[24]

Dictionaries that approximate a given class of signals using few coefficients can be designed using dictionary learning.[9] Dictionaries can be generated ad-hoc from common shape functions such as wavelets or curvelets, however extensive analysis is required to find an optimal set of prescribed shape functions. Dictionary learning proposes a more direct approach: given enough signal examples for a given signal class, learn a dictionary of shape functions that approximate signals within the class using few coefficients. These learned dictionaries (LDs) have improved compression and denoising results for image and video data over ad-hoc dictionaries.[9,11] Dictionary learning has been applied to denoising problems in seismics[25] and ocean acoustics[26,27], as well as to structural acoustic health monitoring.[28]

The K-SVD algorithm,[12] a popular dictionary learning method, finds a dictionary of vectors that optimally partition the data from the training set such that the few dictionary vectors describe each data example. Relative to EOFs which are derived using principal component analysis (PCA),[29,30] these LDs are not constrained to be orthogonal. Thus LDs provide potentially better signal compression because the vectors are on average, nearer to the signal examples (see Fig. 1).[13]

In this paper, LDs describing 1D ocean SSP data from the HF-97 experiment[31,32] and from the South China Sea (SCS)[33] are generated using the K-SVD algorithm and the reconstruction performance is evaluated against EOF methods. In Section II, EOFs, sparse reconstruction methods, and compression are introduced. In Section III, the K-SVD dictionary learning algorithm is explained. In Section IV, SSP reconstruction results are given for LDs and EOFs. It is shown that each shape function within the resulting LDs explain more SSP variability than the leading order EOFs trained on the same

[a] Corresponding author. Electronic mail: mbianco@ucsd.edu

data. Further, it is demonstrated that SSPs can be reconstructed up to acceptable error using as few as one non-zero coefficient. This compression can improve the resolution of ocean SSP estimates with negligible computational burden.

Notation: In the following, vectors are represented by bold lower-case letters and matrices by bold uppercase letters. The ℓ_p-norm of the vector $\mathbf{x} \in \mathbb{R}^N$ is defined as $\|\mathbf{x}\|_p = \left(\sum_{n=1}^{N} \left|x_n\right|^p\right)^{1/p}$. Using similar notation, the ℓ_0-norm is defined as $\|\mathbf{x}\|_0 = \sum_{n=1}^{N} \left|x_n\right|^0 = \sum_{n=1}^{N} 1_{|x_n|>0}$. The ℓ_p-norm of the matrix $\mathbf{A} \in \mathbb{R}^{K\times M}$ is defined as $\|\mathbf{A}\|_p = \left(\sum_{m=1}^{M}\sum_{k=1}^{K} \left|a_k^m\right|^p\right)^{1/p}$. The Frobenius norm ($\ell_2$-norm) of the matrix $\mathbf{A}$ is written as $\|\mathbf{A}\|_{\mathcal{F}}$. The hat symbol $\widehat{\ }$ appearing above vectors and matrices indicates approximations to the true signals or coefficients.

II. EOFS AND COMPRESSION

A. EOFs and PCA

Empirical orthogonal function (EOF) analysis seeks to reduce the dimension of continuously sampled space-time fields by finding spatial patterns which explain much of the variance of the process. These spatial patterns or EOFs correspond to the principal components, from principal component analysis (PCA), of the temporally varying field.[29] Here, the field is a collection of zero-mean ocean SSP anomaly vectors $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_M] \in \mathbb{R}^{K\times M}$, which are sampled over K discrete points in depth and M instants in time. The mean value of the M original observations are subtracted to obtain $\mathbf{Y}$. The variance of the SSP anomaly at each depth sample k, σ_k^2, is defined as

$$\sigma_k^2 = \frac{1}{M}\sum_{m=1}^{M} \left(y_m^k\right)^2 \tag{1}$$

where $[y_1^k, ..., y_M^k]$ are the SSP anomaly values at depth sample k for M time samples.

The singular value decomposition (SVD)[34] finds the EOFs as the eigenvectors of $\mathbf{Y}\mathbf{Y}^{\mathrm{T}}$ by

$$\mathbf{Y}\mathbf{Y}^{\mathrm{T}} = \mathbf{P}\mathbf{\Lambda}^2\mathbf{P}^{\mathrm{T}}, \tag{2}$$

where $\mathbf{P} = [\mathbf{p}_1, ..., \mathbf{p}_L] \in \mathbb{R}^{K\times L}$ are EOFs (eigenvectors) and $\mathbf{\Lambda}^2 = \mathrm{diag}([\lambda_1^2, ..., \lambda_{\mathrm{L}}^2]) \in \mathbb{R}^{\mathrm{L}\times\mathrm{L}}$ are the total variances of the data along the principal directions defined by the EOFs $\mathbf{p}_l$ with

$$\sum_{k=1}^{K}\sigma_k^2 = \frac{1}{M}\mathrm{tr}\left(\mathbf{\Lambda}^2\right). \tag{3}$$

The EOFs $\mathbf{p}_l$ with $\lambda_1^2 \geq ... \geq \lambda_L^2$ are spatial features of the SSPs which explain the greatest variance of $\mathbf{Y}$. If the number of training vectors $M \geq K$, $L = K$ and $[\mathbf{p}_1, ..., \mathbf{p}_L]$ form a basis in $\mathbb{R}^K$.

B. SSP reconstruction using EOFs

Since the leading-order EOFs often explain much of the variance in $\mathbf{Y}$, the representation of anomalies $\mathbf{y}_m$ can be compressed by retaining only the leading order EOFs $P < L$

$$\widehat{\mathbf{y}}_m = \mathbf{Q}_P\widehat{\mathbf{x}}_{P,m} \tag{4}$$

where $\mathbf{Q}_P \in \mathbb{R}^{K\times P}$ is here the dictionary containing the P leading-order EOFs and $\widehat{\mathbf{x}}_{P,m} \in \mathbb{R}^P$ is the coefficient vector. Since the entries in $\mathbf{Q}_P$ are orthonormal, the coefficients are solved by

$$\widehat{\mathbf{x}}_{P,m} = \mathbf{Q}_P^{\mathrm{T}}\mathbf{y}_m. \tag{5}$$

For ocean SSPs, usually no more than $P = 5$ EOF coefficients have been used to reconstruct ocean SSPs.[7,4]

C. Sparse reconstruction

A signal $\mathbf{y}_m$, whose model is sparse in the dictionary $\mathbf{Q}_N = [\mathbf{q}_1, ..., \mathbf{q}_N] \in \mathbb{R}^{K\times N}$ (N-entry sparsifying dictionary for $\mathbf{Y}$), is reconstructed to acceptable error using $T \ll K$ vectors $\mathbf{q}_n$.[9] The problem of estimating few coefficients in $\mathbf{x}_m$ for reconstruction of $\mathbf{y}_m$ can be phrased using the canonical sparse processor

$$\widehat{\mathbf{x}}_m = \underset{\mathbf{x}_m\in\mathbb{R}^N}{\arg\min}\|\mathbf{y}_m - \mathbf{Q}\mathbf{x}_m\|_2 \ \text{ subject to } \ \|\mathbf{x}_m\|_0 \leq T. \tag{6}$$

The ℓ_0-norm penalizes the number of non-zero coefficients in the solution to a typical ℓ_2-norm cost function. The ℓ_0-norm constraint is non-convex and imposes combinatorial search for the exact solution to Eq. (6). Since exhaustive search generally requires a prohibitive number of computations, approximate solution methods such as matching pursuit (MP) and basis pursuit (BP) are preferred.[9] In this paper, orthogonal matching pursuit (OMP)[35] is used as the sparse solver. For small T, OMP achieves similar reconstruction accuracy relative to BP methods, but with much greater speed.[9]

It has been shown that non-orthogonal, overcomplete dictionaries $\mathbf{Q}_N$ with $N > K$ (complete, $N = K$) can be designed to minimize both error and number of non-zero coefficients T, and thus provide greater compression over orthogonal dictionaries.[16,13,9] While overcomplete dictionaries can be designed by concatenating ortho-bases of wavelets or Fourier shape functions, better compression is often achieved by adapting the dictionary to the data under analysis using dictionary learning techniques.[12,13] Since Eq. (6) promotes sparse solutions, it provides criteria for the design of dictionary $\mathbf{Q}$ for adequate reconstruction of $\mathbf{y}_m$ with a minimum number of non-zero coefficients. Rewriting Eq.(7) with

$$\min_{\mathbf{Q}}\left\{\min_{\mathbf{X}}\|\mathbf{Y} - \mathbf{Q}\mathbf{X}\|_{\mathcal{F}}^2 \ \text{ subject to } \ \forall_m, \|\mathbf{x}_m\|_0 \leq T\right\}, \tag{7}$$

where $\mathbf{X} = [\mathbf{x}_1, ..., \mathbf{x}_M]$ is the matrix of coefficient vectors corresponding to examples $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_M]$, reconstruction error is minimized relative to the dictionary $\mathbf{Q}$ as well as relative to the sparse coefficients.

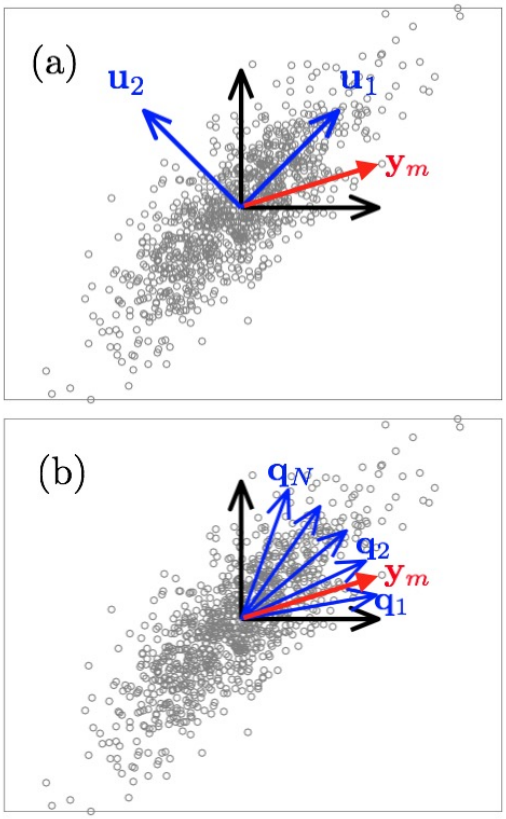

FIG. 1. (Color online) (a) EOF vectors $[\mathbf{u}_1, \mathbf{u}_2]$ and (b) overcomplete LD vectors $[\mathbf{q}_1, ..., \mathbf{q}_N]$ for arbitrary 2D gaussian distribution relative to arbitrary 2D data observation $\mathbf{y}_m$.

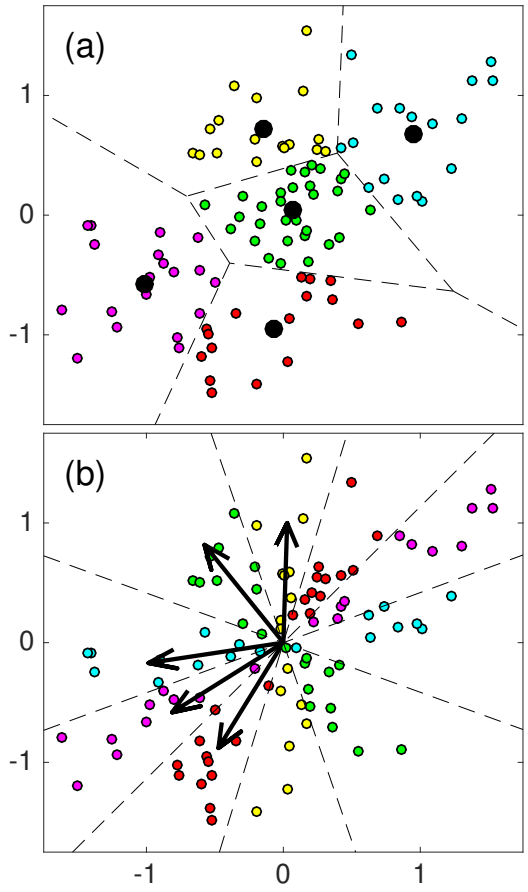

FIG. 2. (Color online) Partitioning of gaussian random distribution ($\sigma_1 = 0.75, \sigma_2 = 0.5$) using (a) 5 codebook vectors (K-means, VQ) and with (b) 5 dictionary vectors from dictionary learning (K-SVD, $T = 1$).

In this paper, the K-SVD algorithm, a clustering based dictionary learning method, is used to solve Eq.(7). The K-SVD is an adaptation of the K-means algorithm for vector quantization (VQ) codebook design (a.k.a. the generalized Lloyd algorithm).[16] The learned dictionary (LD) vectors $\mathbf{q}_n$ from this technique partition the feature space of the data rather than $\mathbb{R}^K$, increasing the likelihood that $\mathbf{y}_m$ is as a linear combination of few vectors $\mathbf{q}_n$ in the solution to Eq. (6) (see Fig. 1). By increasing the number of vectors $N \geq K$ for overcomplete dictionaries, and thus the number of partitions in feature space, the sparsity of the solutions can be increased further.[13]

D. Vector quantization

Vector quantization (VQ)[16] compresses a class of K–dimensional signals $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_M] \in \mathbb{R}^{K \times M}$ by optimally mapping $\mathbf{y}_m$ to a set of code vectors $\mathbf{C} = [\mathbf{c}_1, ..., \mathbf{c}_N] \in \mathbb{R}^{K \times N}$ for $N < M$, called a codebook. The signals $\mathbf{y}_m$ are then quantized or replaced by the best code vector choice from $\mathbf{C}$.[16] The mapping that minimizes mean squared error (MSE) in reconstruction

$$\mathrm{MSE}(\mathbf{Y}, \widehat{\mathbf{Y}}) = \frac{1}{\mathrm{N}} \|\mathbf{Y} - \widehat{\mathbf{Y}}\|_{\mathcal{F}}^2, \tag{8}$$

where $\widehat{\mathbf{Y}} = [\widehat{\mathbf{y}}_1, ..., \widehat{\mathbf{y}}_M]$ is the vector quantized $\mathbf{Y}$, is the assignment of each vector $\mathbf{y}_m$ to the code vectors $\mathbf{c}_n$ based on minimum ℓ_2–distance (nearest neighbor metric). Thus the ℓ_2–distances from the code vectors $\mathbf{c}_n$ define a set of partitions $(R_1, ..., R_N) \in \mathbb{R}^K$ (called Voronoi cells)

$$R_n = \{i \mid \forall_{l \neq n}, \|\mathbf{y}_i - \mathbf{c}_n\|_2 < \|\mathbf{y}_i - \mathbf{c}_l\|_2\}\,, \tag{9}$$

where if $\mathbf{y}_i$ falls within the cell R_n, $\widehat{\mathbf{y}}_i$ is $\mathbf{c}_n$. These cells are shown in Fig. 2(a). This is stated formally by defining a selector function S_n as

$$S_n(\mathbf{y}_m) = \begin{cases} 1 & \text{if } \mathbf{y}_m \in R_n \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

The vector quantization step is then

$$\widehat{\mathbf{y}}_m = \sum_{n=1}^{N} S_n(\mathbf{y}_m)\mathbf{c}_n. \tag{11}$$

The operations in Eq. (9–10) are analogous to solving the sparse minimization problem

$$\widehat{\mathbf{x}}_m = \underset{\mathbf{x}_m \in \mathbb{R}^N}{\arg\min} \|\mathbf{y}_m - \mathbf{C}\mathbf{x}_m\|_2 \;\text{ subject to }\; \|\mathbf{x}_m\|_0 = 1, \tag{12}$$

where the non-zero coefficients $x_m^n = 1$. In this problem, selection of the coefficient in $\mathbf{x}_m$ corresponds to mapping the observation vector $\mathbf{y}_m$ to $\mathbf{c}_n$, similar to the selector function S_n. The vector quantized $\mathbf{y}_m$ is thus written, alternately from Eq. (11), as

$$\widehat{\mathbf{y}}_m = \mathbf{C}\widehat{\mathbf{x}}_m. \tag{13}$$

E. K-means

Given the MSE metric (Eq. (8)), VQ codebook vectors $[\mathbf{c}_1, ..., \mathbf{c}_N]$ which correspond to the centroids of the data $\mathbf{Y}$ within $(R_1, ..., R_N)$ minimize the reconstruction error. The assignment of $\mathbf{c}_n$ as the centroid of $\mathbf{y}_j \in R_n$ is

$$\mathbf{c}_n = \frac{1}{|R_n|} \sum_{j \in R_n} \mathbf{y}_j, \tag{14}$$

where $|R_n|$ is the number of vectors $\mathbf{y}_j \in R_n$.

The K-means algorithm shown in Table I, iteratively updates $\mathbf{C}$ using the centroid condition Eq. (14) and the ℓ_2 nearest–neighbor criteria Eq. (9) to optimize the code vectors for VQ. The algorithm requires an initial codebook $\mathbf{C}^0$. For example, $\mathbf{C}^0$ can be N random vectors in $\mathbb{R}^K$ or selected observations from the training set $\mathbf{Y}$. The K-means algorithm is guaranteed to improve or leave unchanged the MSE distortion after each iteration and converges to a local minimum.[16,12]

TABLE I. The K-means algorithm (Ref. 16.)

	Given: Training vectors $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_M] \in \mathbb{R}^{K \times M}$
	Initialize: index $i = 0$, codebook $\mathbf{C}^0 = [\mathbf{c}_1^0, ..., \mathbf{c}_N^0] \in \mathbb{R}^{K \times N}$, MSE^0 solving Eq. (8)–(11)
I:	Update codebook
	1. Partition $\mathbf{Y}$ into N regions $(R_1, ..., R_N)$ by $R_n = \left\{ i \mid \forall_{l \neq n}, \|\mathbf{y}_i - \mathbf{c}_n^i\|_2 < \|\mathbf{y}_i - \mathbf{c}_l^i\|_2 \right\}$ (Eq. (9))
	2. Make code vectors centroids of $\mathbf{y}_j$ in partitions R_n $\mathbf{c}_n^{i+1} = \frac{1}{\lvert R_n^i \rvert} \sum_{j \in R_n^i} \mathbf{y}_j$
II.	Check error
	1. Calculate MSE^{i+1} from updated codebook $\mathbf{C}^{i+1}$
	2. If $\lvert \mathrm{MSE}^{i+1} - \mathrm{MSE}^{i} \rvert < \eta$ $i = i + 1$, return to I else
	end

III. DICTIONARY LEARNING

Two popular algorithms for dictionary learning, the method of optimal directions (MOD)[13] and the K-SVD,[12] are inspired by the iterative K-means codebook updates for VQ (Table I). The N columns of the dictionary $\mathbf{Q}$, like the entries in codebook $\mathbf{C}$, correspond to partitions in $\mathbb{R}^K$. However, they are constrained to have unit ℓ_2-norm and thus separate the magnitude (coefficients $\mathbf{x}_n$) from the shapes (dictionary entries $\mathbf{q}_n$) for the sparse processing objective Eq.(6). When $T = 1$, the ℓ_2-norm in Eq. (6) is minimized by the dictionary entry $\mathbf{q}_n$ that has the greatest inner product with example $\mathbf{y}_m$.[9] Thus for $T = 1$, $[\mathbf{q}_1, ..., \mathbf{q}_N]$ define radial partitions of $\mathbb{R}^K$. These partitions are shown in Fig. 2(b) for a hypothetical 2D ($K = 2$) random data set. This corresponds to a special case of VQ, called gain-shape VQ.[16] However, for sparse processing, only the shapes of the signals are quantized. The gains, which are the coefficients $\mathbf{x}_m$, are solved. For $T > 1$, the sparse solution is analogous to VQ, assigning examples $\mathbf{y}_m$ to dictionary entries in $\mathbf{Q}$ for up to T non-zero coefficients in $\mathbf{x}_m$.

Given these relationships between sparse processing with dictionaries and VQ, the MOD[13] and K-SVD[12] algorithms attempt to generalize the K-means algorithm to optimization of dictionaries for sparse processing for $T \geq 1$. They are two-step algorithms which reflect the two update steps in the K-means codebook optimization: (1) partition data $\mathbf{Y}$ into regions $(R_1, ..., R_N)$ corresponding to $\mathbf{c}_n$ and (2) update $\mathbf{c}_n$ to centroid of examples $\mathbf{y}_m \in R_N$. The K-means algorithm is generalized to the dictionary learning problem Eq.(7) as two steps:

1. Sparse coding: Given dictionary $\mathbf{Q}$, solve for up to T non-zero coefficients in $\mathbf{x}_m$ corresponding to examples $\mathbf{y}_m$ for $m = [1, ..., M]$

2. Dictionary update: Given coefficients $\mathbf{X}$, solve for $\mathbf{Q}$ which minimizes reconstruction error for $\mathbf{Y}$.

The sparse coding step (1), which is the same for both MOD and K-SVD, is accomplished using any sparse solution method, including matching pursuit and basis pursuit. The algorithms differ in the dictionary update step.

A. The K-SVD Algorithm

The K-SVD algorithm is here chosen for its computational efficiency, speed, and convergence to local minima (at least for $T = 1$). The K-SVD algorithm sequentially optimizes the dictionary entries $\mathbf{q}_n$ and coefficients $\mathbf{x}_m$ for each update step using the SVD, and thus also avoids the matrix inverse. For $T = 1$, the sequential updates of the K-SVD provide optimal dictionary updates for gain-shape VQ.[12,16] Optimal updates to the gain-shape dictionary will, like K-means updates, either improve or leave unchanged the MSE and convergence to a local minimum is guaranteed. For $T > 1$, convergence of the K-SVD updates to a local minimum depends on the accuracy of the sparse-solver used in the sparse coding stage.[12]

In the K-SVD algorithm, each dictionary update step i sequentially improves both the entries $\mathbf{q}_n \in \mathbf{Q}^i$ and the coefficients in $\mathbf{x}_m \in \mathbf{X}^i$, without change in support. Expressing the coefficients as row vectors $\mathbf{x}_T^n \in \mathbb{R}^N$ and $\mathbf{x}_T^j \in \mathbb{R}^N$, which relate all examples $\mathbf{Y}$ to $\mathbf{q}_n$ and $\mathbf{q}_j$,

respectively, the ℓ_2-penalty from Eq. (7) is rewritten as

$$\|\mathbf{Y}-\mathbf{Q}\mathbf{X}\|_{\mathcal{F}}^2 = \left\|\mathbf{Y}-\sum_{n=1}^{N}\mathbf{q}_n\mathbf{x}_T^n\right\|_{\mathcal{F}}^2 = \|\mathbf{E}_j-\mathbf{q}_j\mathbf{x}_T^j\|_{\mathcal{F}}^2, \quad (15)$$

where

$$\mathbf{E}_j = \left(\mathbf{Y}-\sum_{n\neq j}\mathbf{q}_n\mathbf{x}_T^n\right). \quad (16)$$

Thus, in Eq. (15) the ℓ_2-penalty is separated into an error term $\mathbf{E}_j = [\mathbf{e}_{j,1},...,\mathbf{e}_{j,M}] \in \mathbb{R}^{K\times M}$, which is the error for all examples $\mathbf{Y}$ if $\mathbf{q}_j$ is excluded from their reconstruction, and the product of the excluded entry $\mathbf{q}_j$ and coefficients $\mathbf{x}_T^j \in \mathbb{R}^N$.

An update to the dictionary entry $\mathbf{q}_j$ and coefficients $\mathbf{x}_T^j$ which minimizes Eq. (15) is found by taking the SVD of $\mathbf{E}_j$, which provides the best rank-1 approximation of $\mathbf{E}_j$. However, many of the entries in $\mathbf{x}_T^j$ are zero (corresponding to examples which don't use $\mathbf{q}_j$). To properly update $\mathbf{q}_j$ and $\mathbf{x}_T^j$ with SVD, Eq. (15) must be restricted to examples $\mathbf{y}_m$ which use $\mathbf{q}_j$

$$\|\mathbf{E}_j^R-\mathbf{q}_j\mathbf{x}_R^j\|_{\mathcal{F}}^2, \quad (17)$$

where $\mathbf{E}_j^R$ and $\mathbf{x}_R^j$ are entries in $\mathbf{E}_j$ and $\mathbf{x}_T^j$, respectively, corresponding to examples $\mathbf{y}_m$ which use $\mathbf{q}_j$, and are defined as

$$\mathbf{E}_j^R = \left\{\mathbf{e}_{j,l}\middle|\forall_l,\ x_l^j \neq 0\right\},\ \mathbf{x}_R^j = \left\{x_l^j\middle|\ \forall_l,\ x_l^j \neq 0\right\}. \quad (18)$$

Thus for each K-SVD iteration, the dictionary entries and coefficients are sequentially updated as the SVD of $\mathbf{E}_j^R = \mathbf{U}\mathbf{S}\mathbf{V}^{\mathrm{T}}$. The dictionary entry $\mathbf{q}_j^i$ is updated with the first column in $\mathbf{U}$ and the coefficient vector $\mathbf{x}_R^j$ is updated as the product of the first singular value $\mathbf{S}(1,1)$ with the first column of $\mathbf{V}$. The K-SVD algorithm is given in Table II.

The dictionary $\mathbf{Q}$ is initialized using N randomly selected, ℓ_2-normalized examples from $\mathbf{Y}$.[12,9] During the iterations, one or more dictionary entries may become unused. If this occurs, the unused entries are replaced using the most poorly represented examples $\mathbf{y}_m$ (ℓ_2-normlized), determined by reconstruction error.

TABLE II. The K-SVD Algorithm (Ref. 12)

	Given: $\mathbf{Y} \in \mathbb{R}^{K\times M}$, $\mathbf{Q}^0 \in \mathbb{R}^{K\times N}$, $T \in \mathbb{N}$, and $i = 0$
	Repeat until convergence:
1.	Sparse coding
	for $m = 1 : M$
	solve Eq. (6) using any sparse solver
a:	$\widehat{\mathbf{x}}_m = \underset{\mathbf{x}_m\in\mathbb{R}^N}{\arg\min}\|\mathbf{y}_m-\mathbf{Q}^i\mathbf{x}_m\|_2$ subject to $\|\mathbf{x}_m\|_0 \leq T$
	end
b:	$\mathbf{X} = [\widehat{\mathbf{x}}_1,...,\widehat{\mathbf{x}}_M]$
2.	Dictionary update
	for $j = 1 : N$
a:	compute reconstruction error $\mathbf{E}_j$ as $\mathbf{E}_j = \mathbf{Y} - \sum_{n\neq j}\mathbf{q}_n^i\mathbf{x}_T^n$
b:	obtain $\mathbf{E}_j^R$, $\mathbf{x}_R^j$ corresponding to nonzero $\mathbf{x}_T^j$
c:	apply SVD to $\mathbf{E}_j^R$ $\mathbf{E}_j^R = \mathbf{U}\mathbf{S}\mathbf{V}^{\mathrm{T}}$
d:	update $\mathbf{q}_j^i$: $\mathbf{q}_j^i = \mathbf{U}(:,1)$
e:	update $\mathbf{x}_R^j$: $\mathbf{x}_R^j = \mathbf{V}(:,1)\mathbf{S}(1,1)$
	end
f:	$\mathbf{Q}^{i+1} = \mathbf{Q}^i$
	$i = i+1$

IV. EXPERIMENTAL RESULTS

To demonstrate the usefulness of the dictionary learning approach, we here analyze two data sets: (1) thermistor data from the HF-97 acoustics experiment,[31,32] conducted off the coast of Point Loma, CA and (2) conductivity, temperature, and depth (CTD) data collected across the Luzon Strait near the South China Sea (SCS).[33] Training data $\mathbf{Y}$ were derived from the data sets by converting raw thermistor and CTD data to SSPs and subtracting the mean. The HF-97 thermistor data was recorded every 15 s, over a 48 hour period, from 14 to 70 m depth, with 4 m spacing (15 points). The full 11,488 profile data set was down-sampled to $M = 1000$ profiles for the training set, and SSPs were interpolated to $K = 30$ points using a shape-preserving cubic spline. The SCS CTD data was recorded at about 1 m resolution from 116 to 496 m depth (384 points). From the SCS data set, $M = 755$ profiles were used as the training set, and the profiles were uniformly down-sampled to $K = 50$ points. The SSP data sets are shown in Fig. 3. Both data sets have small and large spatiotemporal variations.

EOFs were calculated from the SVD (Eq. 2) and LDs (learned dictionaries) were generated with the K-SVD algorithm (Table II), using OMP for the sparse coding stage. The number of non-zero coefficients solved with OMP for each dictionary was held fixed at exactly T nonzero coefficients. The initial dictionary $\mathbf{Q}^0$ was populated using randomly selected examples from the training sets $\mathbf{Y}$.

A. Learning SSP dictionaries from data

Here, LDs and EOFs were generated using the full SSP data from HF-97 ($M = 1000$) and SCS ($M = 755$). The EOFs and LDs from HF-97 are shown in Fig. 4–5 and from the SCS in Fig. 6. The HF-97 LD, with $N = K$ and $T = 1$, is compared to the EOFs ($K = 30$) in Fig. 4. Only the leading order EOFs (Fig. 4(a)) are informative of ocean SSP variability whereas all shape functions in the LD (Fig. 4(b)) are informative (Fig. 4(c)–(d)). This behavior is also evident for the SCS data set (Fig. 6). The EOFs ($K = 50$) calculated from the full training set

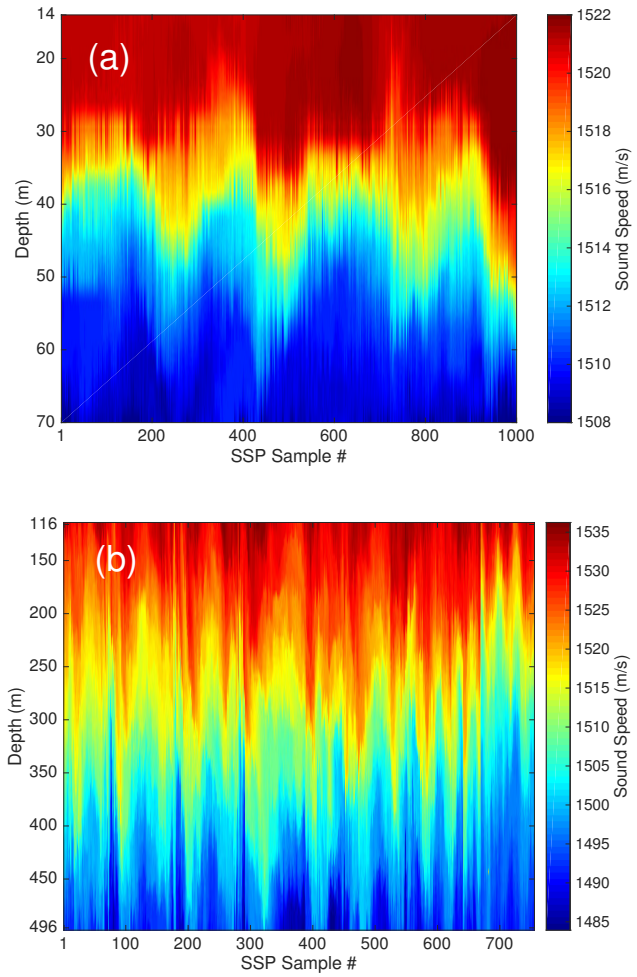

FIG. 3. (Color online) SSP data from (a) HF-97 experiment and (b) SCS.

are shown in Fig. 6(a), and the LD entries for $N = 50$ and $T = 1$ sparse coefficient are shown in Fig. 6(b). The overcomplete LDs for the HF-97 data shown in Fig. 5 and for the SCS data in Fig. 6(c).

As illustrated in Fig. 1, by relaxing the requirement of orthogonality for the shape functions, the shape functions can better fit the data and thereby achieve greater compression. The Gram matrix $\mathbf{G}$, which gives the coherence of matrix columns, is defined for a matrix $\mathbf{A}$ with unit ℓ_2-norm columns as $\mathbf{G} = |\mathbf{A}^{\mathrm{T}}\mathbf{A}|$. The Gram matrix for the EOFs (Fig. 4(e)) shows the shapes in the EOF dictionary are orthogonal ($\mathbf{G} = \mathbf{I}$, by definition), whereas those of the LD (Fig. 4(f)) are not.

B. Reconstruction of SSP training data

In this section, EOFs and LDs are trained on the full SSP data sets $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_M]$. Reconstruction performance of the EOF and LDs are then evaluated on SSPs within the training set, using a mean error metric.

The coefficients for the learned $\mathbf{Q}$ and initial $\mathbf{Q}^0$ dictionaries $\widehat{\mathbf{x}}_m$ are solved from the sparse objective (Eq. (6)) using OMP. The least squares (LS) solution for the T leading-order coefficients $\mathbf{x}_L \in \mathbb{R}^T$ from the EOFs $\mathbf{P}$ were solved by Eq. (5). The best combination of T EOF coefficients was solved from the sparse objective (Eq. (6)) using OMP. Given the coefficients $\mathbf{X} = [\mathbf{x}_1, ..., \mathbf{x}_m]$ describing examples $\mathbf{Y} = [\mathbf{y}_1, ..., \mathbf{y}_m]$, the reconstructed examples $\widehat{\mathbf{Y}} = [\widehat{\mathbf{y}}_1, ..., \widehat{\mathbf{y}}_m]$ are given by $\widehat{\mathbf{Y}} = \mathbf{Q}\widehat{\mathbf{X}}$. The mean reconstruction error ME for the training set is then

$$\mathrm{ME} = \frac{1}{KM}\|\mathbf{Y} - \widehat{\mathbf{Y}}\|_1. \tag{19}$$

We here use the ℓ_1-norm to stress the robustness of the LD reconstruction.

To illustrate the optimality of LDs for SSP compression, the K-SVD algorithm was run using EOFs as the initial dictionary $\mathbf{Q}^0$ for $T = 1$ non-zero coefficient. The convergence of ME for the K-SVD iterations is shown in Fig. 7(a). After 30 K-SVD iterations, the mean error of the $M = 1000$ profile training set is decreased by nearly half. The convergence is much faster for $\mathbf{Q}^0$ consisting of randomly selected examples from $\mathbf{Y}$.

For LDs, increasing the number of entries N or increasing the number of sparse coefficients T will always reduce the reconstruction error (N and T are decided with computational considerations). The effect of N and T on the mean reconstruction error for the HF-97 data is shown in Fig. 7(b). The errors are calculated for the range $N = K$ to $N = 4K$ and the dictionaries were optimized to use a fixed number non-zero coefficients (T).

The reconstruction error using the EOF dictionary is compared to results from LDs $\mathbf{Q}$ with $N = 3K$, using T non-zero coefficients. In Fig. 8[(a) and (c)] results are shown for the HF-97 ($N = 90$) and SCS ($N = 150$) data, respectively. Coefficients describing each example $\mathbf{y}_m$, were solved (1) from the LD $\mathbf{Q}$, (2) from $\mathbf{Q}^0$, the dictionary consisting of N randomly chosen examples from the training set (to illustrate improvements in reconstruction error made in the K-SVD iterations), (3) the leading order EOFs, and (4) the best combination of EOFs. The mean SSP reconstruction error using the LDs trained for each sparsity T is less than EOF reconstruction, for either leading order coefficients or best coefficient combination, for all values of T shown. The best combination of EOF coefficients, chosen approximately using OMP, achieves less error than the LS solution to the leading order EOFs, with added cost of search.

Just one LD entry achieves the same ME as more than 6 leading order EOF coefficients, or greater than 4 EOF coefficients chosen by search (Fig. 8[(a) and (c)]). To illustrate the representational power of the LD entries, both true and reconstructed SSPs are shown in Fig. 9(a) for the HF-97 data and in Fig. 9(b) for the SCS data. Nine true SSP examples from each training set, for HF-97 (SCS) taken at 100 (80) point intervals from $m = 100$ to 900 (80 to 720), are reconstructed using one LD coefficient. It is shown for each case, that nearly all of the SSP variability is captured using a single LD coefficient.

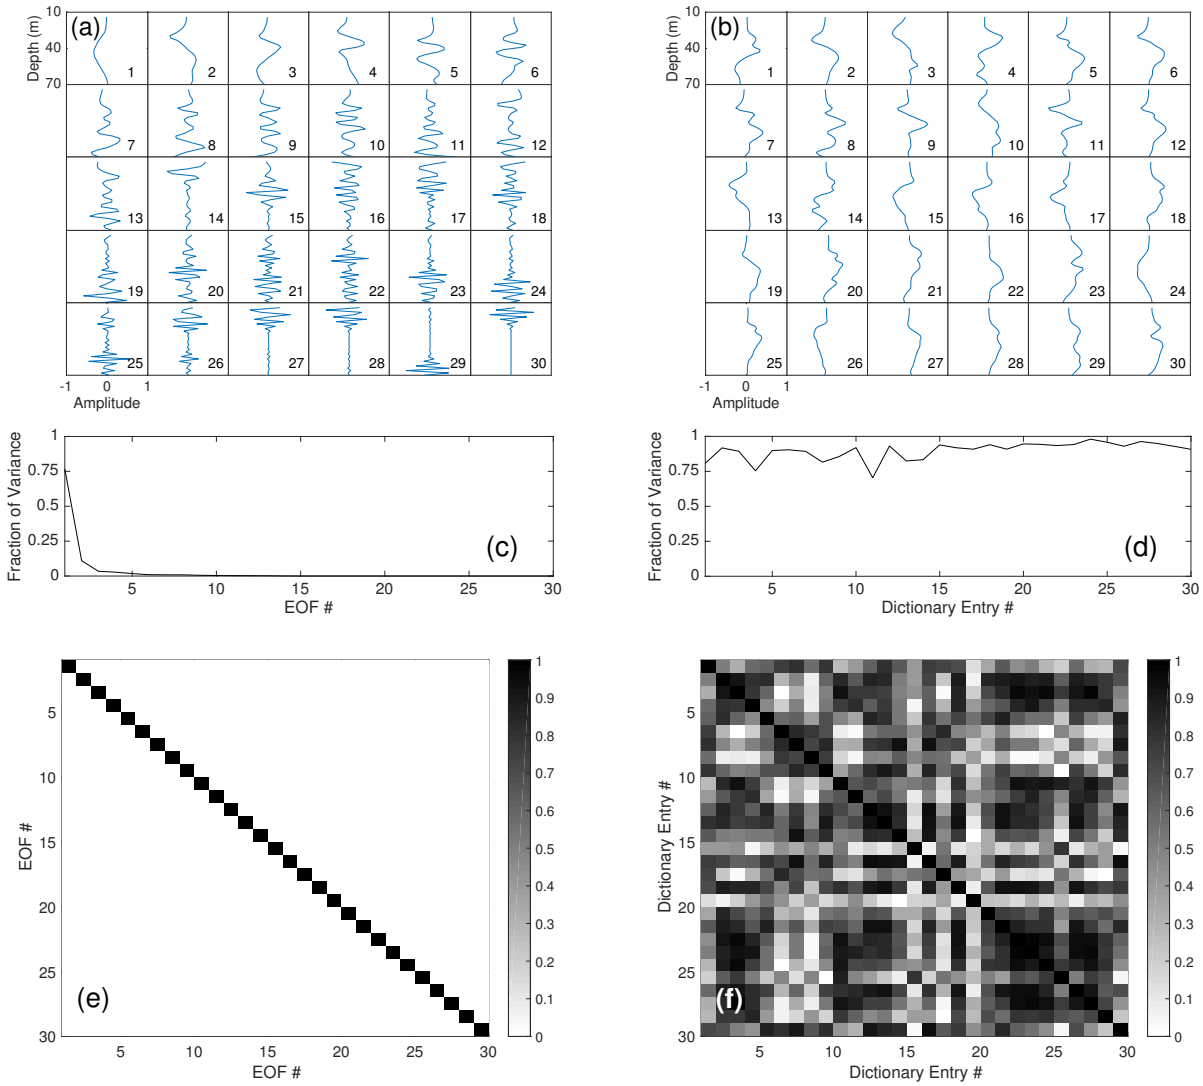

FIG. 4. (Color online) HF-97: (a) EOFs and (b) LD entries ($N = K$ and $T = 1$, sorted by variance $\sigma^2_{\mathbf{q}_n}$). Fraction of (c) total SSP variance explained by EOFs and (d) SSP variance explained for examples using LD entries. Coherence of (e) EOFs and (f) LD entries.

C. Cross-validation of SSP reconstruction

The out of sample SSP reconstruction performance of LDs and EOFs is tested using K-fold cross-validation.[34] The entire SSP data set $\mathbf{Y}$ of M profiles, for each experiment, is divided into J subsets with equal numbers of profiles $\mathbf{Y} = [\mathbf{Y}_1, ..., \mathbf{Y}_J]$, where the fold $\mathbf{Y}_j \in \mathbb{R}^{K\times(M/J)}$. For each of the J folds: (1) $\mathbf{Y}_j$ is the set of out of sample test cases, and the training set $\mathbf{Y}_{tr}$ is

$$\mathbf{Y}_{tr} = \left\{\mathbf{Y}_l \middle|\ \forall_{l\neq j}\right\}; \tag{20}$$

(2) the LD $\mathbf{Q}_j$ and EOFs are derived using $\mathbf{Y}_{tr}$; and (3) coefficients estimating test samples $\mathbf{Y}_j$ are solved for $\mathbf{Q}_j$ with sparse processor Eq. (6), and for EOFs by solving for leading order terms and by solving with sparse processor. The out of sample error from cross validation $\mathrm{ME_{CV}}$ for each method is then

$$\mathrm{ME}_{CV} = \frac{1}{KM}\sum_{j=1}^{J} \|\mathbf{Y}_j - \widehat{\mathbf{Y}}_j\|_1. \tag{21}$$

The out of sample reconstruction error $\mathrm{ME_{CV}}$ increases over the within-training-set estimates for both the learned and EOF dictionaries, as shown in Fig. 8[(b) and (d)] for $J = 10$ folds. The mean reconstruction error using the LDs, as in the within-training-set estimates, is less than the EOF dictionaries. For both the HF-97 (SCS) data, more than 2 (2) EOF coefficients, choosing best combination by search, or more than 3 (equal to 3) leading-order EOF coefficients solved with LS, are required to achieve the same out of sample performance as one LD entry.

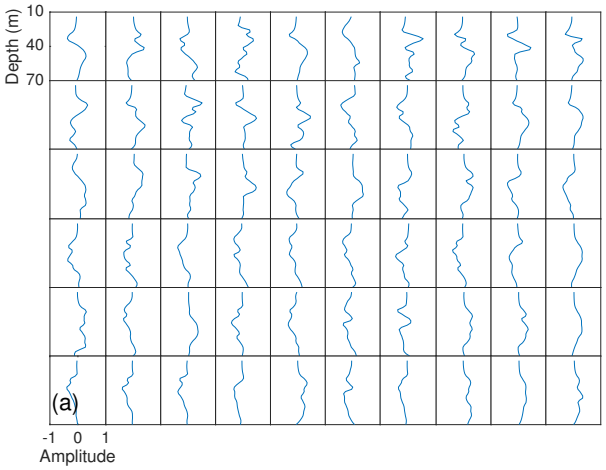

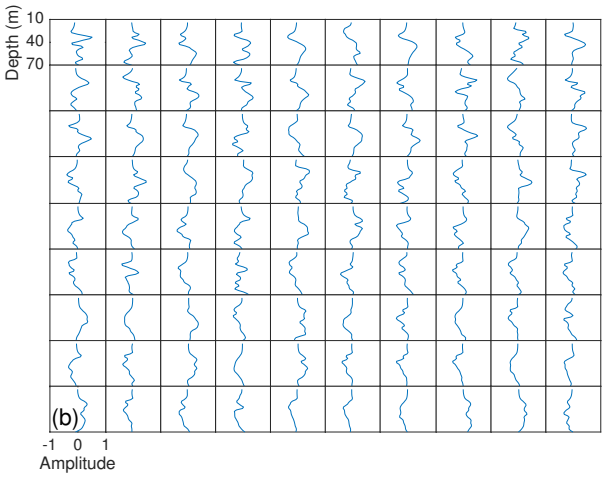

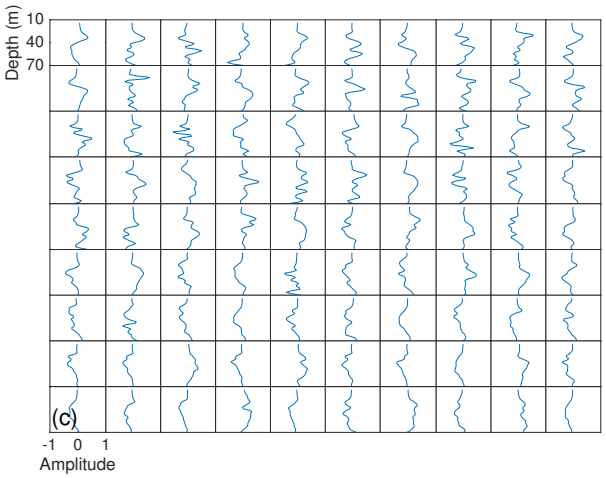

FIG. 5. (Color online) HF-97: LD entries (a) $N = 60$ and $T = 1$, (b) $N = 90$ and $T = 1$, and (c) $N = 90$ and $T = 5$. Dictionary entries are sorted in descending variance $\sigma^2_{\mathbf{q}_n}$.

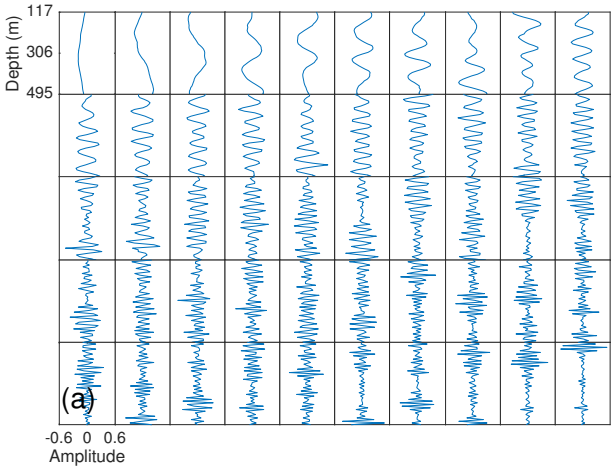

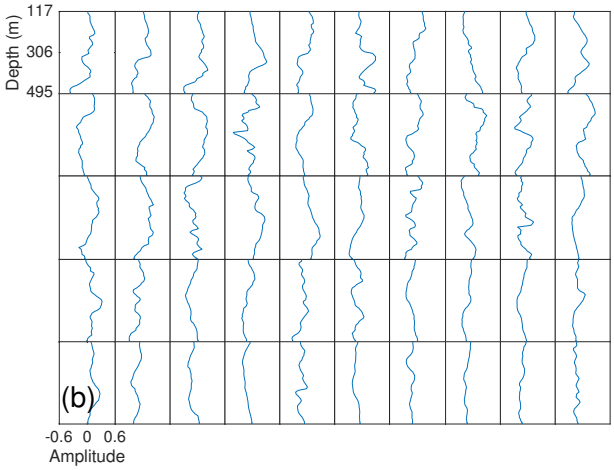

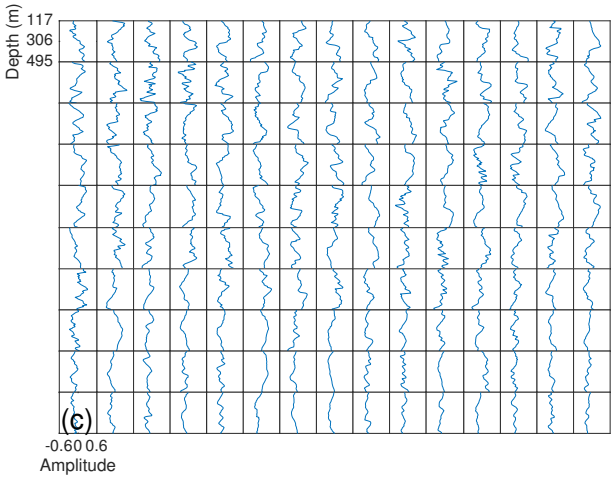

FIG. 6. (Color online) SCS: EOFs (a) and LD entries; (b) $N = K = 50$ and $T = 1$ and (c) $N = 150$ and $T = 1$. Dictionary entries are sorted in descending variance $\sigma^2_{\mathbf{q}_n}$.

D. Solution space for SSP inversion

Acoustic inversion for ocean SSP is a non-linear problem. One approach is coefficient search using genetic algorithms.[1] Discretizing each coefficient into H values, the number of candidate solutions for T fixed coefficients indices is

$$S_{\text{fixed}} = H^T. \tag{22}$$

If the coefficient indices for the solution can vary, as per dictionary learning with LD $\mathbf{Q} \in \mathbb{R}^{K \times N}$, the number of

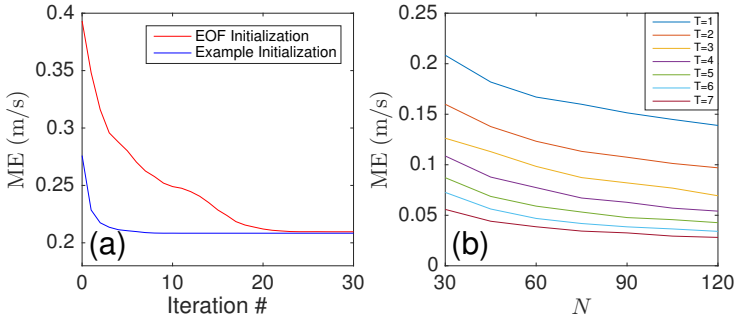

FIG. 7. (Color online) HF-97: (a) Convergence of LD ($N = 30$, $T = 1$) mean reconstruction error (ME), initialized using EOFs or N randomly selected examples from $\mathbf{Y}$. (b) ME versus non-zero coefficients T and number of dictionary entries N.

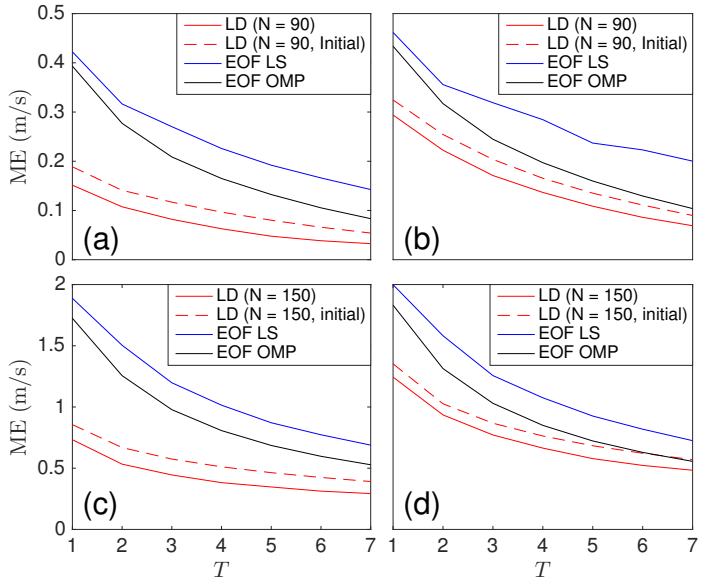

FIG. 8. (Color online) Mean reconstruction error (ME) versus T using EOFs (solved using LS and OMP) and LDs ($N = 90$ for HF-97 and $N = 150$ for SCS) for (a) HF- 97 and (c) SCS. Mean reconstruction error $\mathrm{ME_{CV}}$ for out-of-sample data calculated with K-fold cross validation for J = 10 folds, (b) HF-97 and (d) SCS.

candidate solutions S_{comb} is

$$S_{\mathrm{comb}} = H^T \frac{N!}{T!(N-T)!}. \tag{23}$$

Using a typical $H = 100$ point discretization of the coefficients, the number of possible solutions for fixed and combinatorial dictionary indices are plotted in Fig. 10. Assuming an unknown SSP similar to the training set, the SSP may be constructed up to acceptable resolution using one coefficient from the LD (10^4 possible solutions, see Fig. 10). To achieve the similar ME, 7 EOFs coefficients are required (10^{14} possible solutions, Fig. 10) using fixed indices and the best EOF combination requires 5 EOFs (10^{17} possible solutions, Fig. 10).

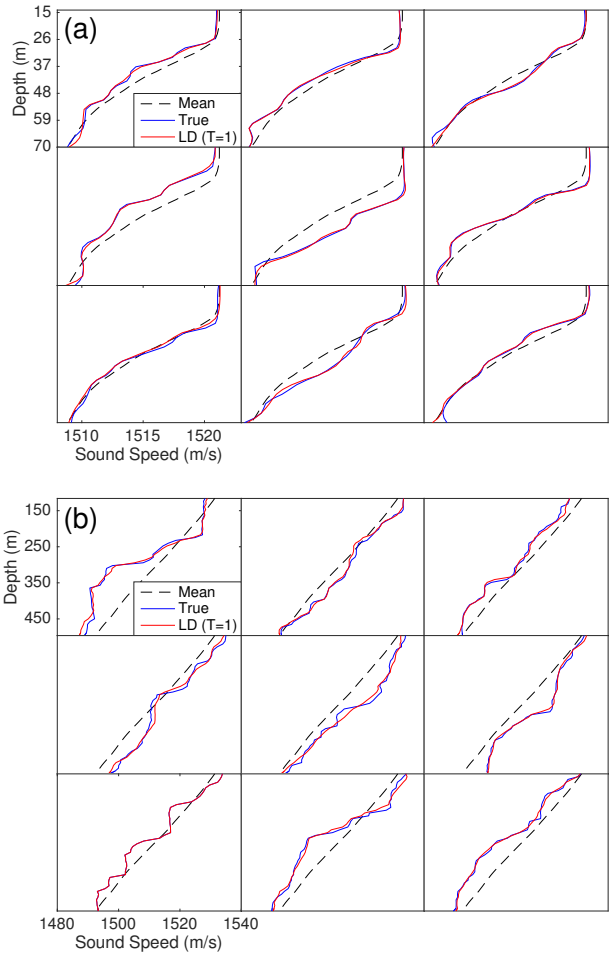

FIG. 9. (Color online) True SSP reconstruction of 9 example profiles using one coefficient ($T = 1$) from LD for (a) HF-97 ($N = 90$) and (b) SCS ($N = 150$).

V. CONCLUSION

Given sufficient training data, dictionary learning generates optimal dictionaries for sparse reconstruction of a given signal class. Since these LDs are not constrained to be orthogonal, the entries fit the distribution of the data such that signal example is approximated using few LD entries. Relative to EOFs, each LD entry is informative to the signal variability.

The K-SVD dictionary learning algorithm is applied to ocean SSP data from the HF-97 and SCS experiments. It is shown that the LDs generated describe ocean SSP variability with high resolution using fewer coefficients than EOFs. As few as one coefficient from a LD describes nearly all the variability in each of the observed ocean SSPs. This performance gain is achieved by the larger number of informative elements in the LDs over EOF dictionaries. Provided sufficient SSP training data is available, LDs can improve SSP inversion resolution with negligible computational expense. This could provide improvements to geoacoustic inversion,[1] matched field processing,[36,37] and underwater communications.[31]

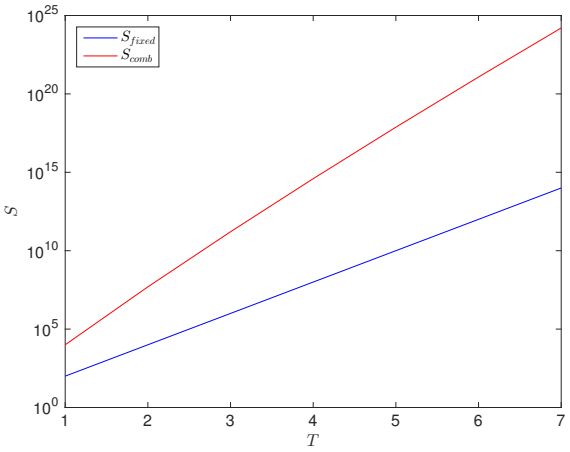

FIG. 10. (Color online) Number of candidate solutions S for SSP inversion versus T, S_{fixed} using fixed indices and S_{comb} best combination of coefficients. Each coefficient is discretized with $H = 100$ for dictionary $\mathbf{Q} \in \mathbb{R}^{K \times N}$ with $N = 100$.

Acknowledgments

The authors would like to thank Dr. Robert Pinkel for the use of the South China Sea CTD data. This work is supported by the Office of Naval Research, Grant No. N00014-11-1-0439.

[1] P. Gerstoft, "Inversion of seismoacoustic data using genetic algorithms and *a posteriori* probability distributions," J. Acoust. Soc. Am. **95**(2), 770–782 (1994).

[2] L.R. LeBlanc and F.H. Middleton. "An underwater acoustic sound velocity data model," J. Acoust. Soc. Am. **67**(6), 2055–2062 (1980).

[3] M. I. Taroudakis and J. S. Papadakis, "A modal inversion scheme for ocean acoustic tomography," J. Comp. Acoust. **1**(4), 395–421 (1993).

[4] P. Gerstoft and D.F. Gingras, "Parameter estimation using multifrequency range–dependent acoustic data in shallow water,"J. Acoust. Soc. Am. **99**(5), 2839–2850 (1996).

[5] C. Park, W. Seong, P. Gerstoft, and W. S. Hodgkiss,"Geoacoustic inversion using backpropagation," IEEE J. Ocean. Eng. **35**(4), 722–731 (2010).

[6] B. A. Tan, P. Gerstoft, C. Yardim, and W. S. Hodgkiss,"Broadband synthetic aperture geoacoustic inversion," J. Acoust. Soc. Am. **134**(1), 312–322 (2013).

[7] C. F. Huang, P. Gerstoft, and W. S. Hodgkiss, "Effect of ocean sound speed uncertainty on matched-field geoacoustic inversion," J. Acoust. Soc. Am. **123**(6), EL162–EL168 (2008).

[8] R. Rubinstein, A.M. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," Proceedings of the IEEE **98**(6), 1045–1057 (2010).

[9] M. Elad, *Sparse and Redundant Representations*, Springer, New York (2010).

[10] I. Tosic and P. Frossard, "Dictionary learning," IEEE Sig. Proc. Mag. **28**(2), 27–38 (2011).

[11] K. Schnass, "On the identifiability of overcomplete dictionaries via the minimisation principle underlying K-SVD," App. and Comp. Harm. Anal. **37**(3), 464–491 (2014).

[12] M. Aharon, M. Elad, and A. Bruckstein "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," IEEE Trans. Sig. Proc. **54**(11), 4311–4322 (2006).

[13] K. Engan, S.O. Aase, J.H. Husøy, "Multi-frame compression: theory and design," Signal Processing **80**(10), 2121–2140 (2000).

[14] A. Hyvärinen, J. Hurri, and P.O. Hoyer, *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*, Springer science and business media (2009).

[15] C. Christopoulos, A. Skodras, and T. Ebrahimi, "The JPEG2000 still image coding system: an overview," IEEE Trans. Cons. Elec. **46**(4), 1103–1127 (2000).

[16] A. Gersho and R.M. Gray, *Vector quantization and signal compression*. Norwell, MA: Kluwer Academic, 1991.

[17] H.L. Taylor, S.C. Banks, and J.F. McCoy, "Deconvolution with the ℓ_1–norm," Geophysics **44**(1), 39–52 (1979).

[18] E. Candés, "Compressive sampling," *Proceedings of the international congress of mathematicians*, Vol. 3, 1433–1452 (2006).

[19] G. Edelmann and C. Gaumond, "Beamforming using compressive sensing," J. Acoust. Soc. Am. **130**(4), EL232-EL237 (2011).

[20] A. Xenaki, P. Gerstoft, and K. Mosegaard, "Compressive beamforming," J. Acoust. Soc. Am. **136**(1), 260–271 (2014).

[21] P. Gerstoft, A. Xenaki, and C. F. Mecklenbräuker, "Multiple and single snapshot compressive beamforming," J. Acoust. Soc. Am. **138**(4), 2003–2014 (2015).

[22] Y. Choo and W. Song, "Compressive spherical beamforming for localization of incipient tip vortex cavitation," J. Acoust. Soc. Am. **140**(6), 4085–4090 (2016).

[23] C. Yardim, P. Gerstoft, W. S. Hodgkiss, and J. Traer, "Compressive geoacoustic inversion using ambient noise," J. Acoust. Soc. Am. **135**(3), 1245–1255 (2014).

[24] M. Bianco and P. Gerstoft, "Compressive acoustic sound speed profile estimation," J. Acoust. Soc. Am. EL **139**(3), EL90–EL94 (2016).

[25] S. Beckouche and J. Ma, "Simultaneous dictionary learning and denoising for seismic data," Geophysics **79**(3), 27–37 (2014).

[26] M. Taroudakis and C. Smaragdakis, "De-noising procedures for inverting underwater acoustic signals in applications of acoustical oceanography," *Euronoise 2015 Maastricht*, pp.1393–1398 (2015).

[27] T. Wang and W. Xu, "Sparsity-based approach for ocean acoustic tomography using learned dictionaries," *OCEANS 2016 Shanghai IEEE*, pp.1–6 (2016).

[28] K.S. Alguri and J.B. Harley, "Consolidating guided wave simulations and experimental data: a dictionary leaning approach," Proc. SPIE, Health Monitoring of Structural and Biological Systems **9805**, 98050Y-1–98050Y-10 (2016).

[29] A. Hannachi, I.T. Jolliffe, and D.B. Stephenson, "Empirical orthogonal functions and related techniques in atmospheric science: a review," International Journal of Climatology **27**(9), 1119–1152 (2007).

[30] A.H. Monahan, J.C. Fyfe, M.H. Ambaum, D.B. Stephenson, and G.R. North, "Empirical orthogonal functions: the medium is the message," Journal of Climate **22**(24), 6501–6514 (2009).

[31] N. Carbone and W.S. Hodgkiss, "Effects of tidally driven temperature fluctuations on shallow-water acoustic communications at 18kHz," IEEE Journal of Ocean. Eng. **25**(1), 84–94 (2000).

[32] W.S. Hodgkiss, W.A. Kuperman, and D.E. Ensberg, "Channel impulse response fluctuations at 6 kHz in shal-

low water," *Impact of littoral environmental variability of acoustic predictions and sonar performance*, Springer, Netherlands, pp. 295–302 (2002).

[33] C.T. Liu, R. Pinkel, M.K. Hsu, J. M. Klymak, H.W. Chen, and C. Villanoy, "Nonlinear internal waves from the Luzon Strait," Eos Trans. AGU **87**(42), 449–451 (2006).

[34] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference and prediction*, 2nd Ed. Springer (2009).

[35] Y.C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in IEEE Proc. 27th. Annu. Asilomar Conf. Signals, Systems and Computers, 40–44 (1993).

[36] A.B. Baggeroer, W.A. Kuperman, and P.N. Mikhalevsky, "An overview of matched field methods in ocean acoustics," IEEE Journal of Ocean. Eng. **18**(4), 401–424 (1993).

[37] C.M. Verlinden, J. Sarkar, W.S. Hodgkiss, W.A. Kuperman, and K.G. Sabra, "Passive acoustic source localization using sources of opportunity," J. Acoust. Soc. Am. **138**(1), EL54–EL59 (2015).